



US009471348B2

(12) **United States Patent**  
**Zuo et al.**

(10) **Patent No.:** **US 9,471,348 B2**  
(45) **Date of Patent:** **Oct. 18, 2016**

(54) **APPLYING POLICIES TO SCHEDULE NETWORK BANDWIDTH AMONG VIRTUAL MACHINES**

USPC ..... 370/230, 235, 235.1, 237, 254–255, 370/395.4; 709/201, 225, 226, 248  
See application file for complete search history.

(71) Applicant: **MICROSOFT TECHNOLOGY LICENSING, LLC**, Remond, WA (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,436,844 B2 10/2008 Wang et al.  
7,447,155 B2 11/2008 Tang

(Continued)

FOREIGN PATENT DOCUMENTS

CN 1663195 A 8/2005  
TW 200618543 A1 6/2006

OTHER PUBLICATIONS

“First Office Action and Search Report Issued for Chinese Patent Application No. 201180026749.8”, Mailed Date: Jan. 7, 2015, 12 Pages.

(Continued)

*Primary Examiner* — Warner Wong

(74) *Attorney, Agent, or Firm* — Ben Tabor; Kate Drakos; Micky Minhas

(57) **ABSTRACT**

Computerized methods, systems, and computer-storage media for allowing virtual machines (VMs) residing on a common physical node to fairly share network bandwidth are provided. Restrictions on resource consumption are implemented to ameliorate stressing the network bandwidth or adversely affecting the quality of service (QoS) guaranteed to tenants of the physical node. The restrictions involves providing a scheduler that dynamically controls networking bandwidth allocated to each of the VMs as a function of QoS policies. These QoS policies are enforced by controlling a volume of traffic being sent from the VMs. Controlling traffic includes depositing tokens into token-bucket queues assigned to the VMs, respectively. The tokens are consumed as packets pass through the token-bucket queues. Upon consumption, packets are held until sufficient tokens are reloaded to the token-bucket queues.

**14 Claims, 9 Drawing Sheets**

(21) Appl. No.: **13/932,952**

(22) Filed: **Jul. 1, 2013**

(65) **Prior Publication Data**

US 2013/0298123 A1 Nov. 7, 2013

**Related U.S. Application Data**

(63) Continuation of application No. 12/790,981, filed on May 31, 2010, now Pat. No. 8,477,610.

(51) **Int. Cl.**

**G06F 9/455** (2006.01)  
**H04L 29/06** (2006.01)

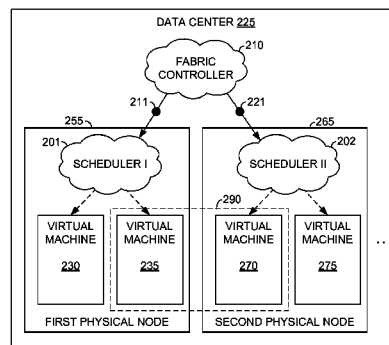
(Continued)

(52) **U.S. Cl.**

CPC ..... **G06F 9/45533** (2013.01); **H04L 29/06523** (2013.01); **H04L 41/0896** (2013.01); **H04L 47/20** (2013.01); **H04L 47/215** (2013.01)

(58) **Field of Classification Search**

CPC ..... G06F 9/45533; H04L 29/06523; H04L 12/2439; H04L 12/40065; H04L 41/0896; H04L 47/525; H04L 47/783; H04L 47/787



## Page 2

|              |      |        |                         |         |
|--------------|------|--------|-------------------------|---------|
| 8,495,633    | B2 * | 7/2013 | Easton et al. ....      | 718/1   |
| 8,516,137    | B2   | 8/2013 | Calder et al. ....      | 709/229 |
| 8,705,351    | B1 * | 4/2014 | McGlaughlin et al. .... | 370/230 |
| 8,782,137    | B2 * | 7/2014 | Jung et al. ....        | 709/204 |
| 2002/0101820 | A1   | 8/2002 | Gupta et al. ....       |         |
| 2006/0184349 | A1 * | 8/2006 | Goud et al. ....        | 703/24  |
| 2008/0002683 | A1 * | 1/2008 | Droux et al. ....       | 370/389 |
| 2008/0002733 | A1 * | 1/2008 | Sutskover ....          | 370/436 |

|              |      |         |                      |         |
|--------------|------|---------|----------------------|---------|
| 2008/0008095 | A1 * | 1/2008  | Gilfix .....         | 370/235 |
| 2009/0238072 | A1 * | 9/2009  | Tripathi et al. .... | 370/235 |
| 2011/0185064 | A1 * | 7/2011  | Head et al. ....     | 709/226 |
| 2011/0213686 | A1 * | 9/2011  | Ferris et al. ....   | 705/34  |
| 2011/0296052 | A1 * | 12/2011 | Guo et al. ....      | 709/240 |
| 2011/0296231 | A1 * | 12/2011 | Dake .....           | 714/42  |

“Second Office Action and Search Report Issued in Chinese Patent Application No. 201180026749.8”, Mailed Date: Jun. 26, 2015, 11 Pages.

Office Action and Search Report Issued in Taiwan Patent Application No. 100112856, Mailed Date: Sep. 2, 2015, 13 Pages.

\* cited by examiner

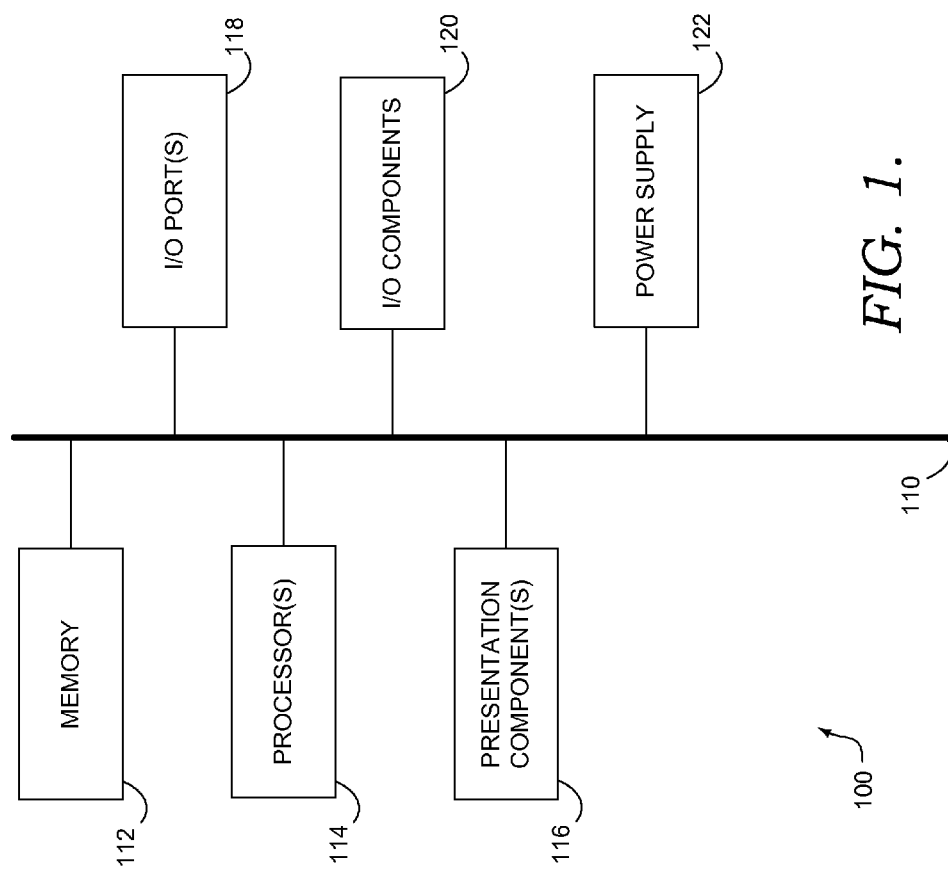


FIG. 1.

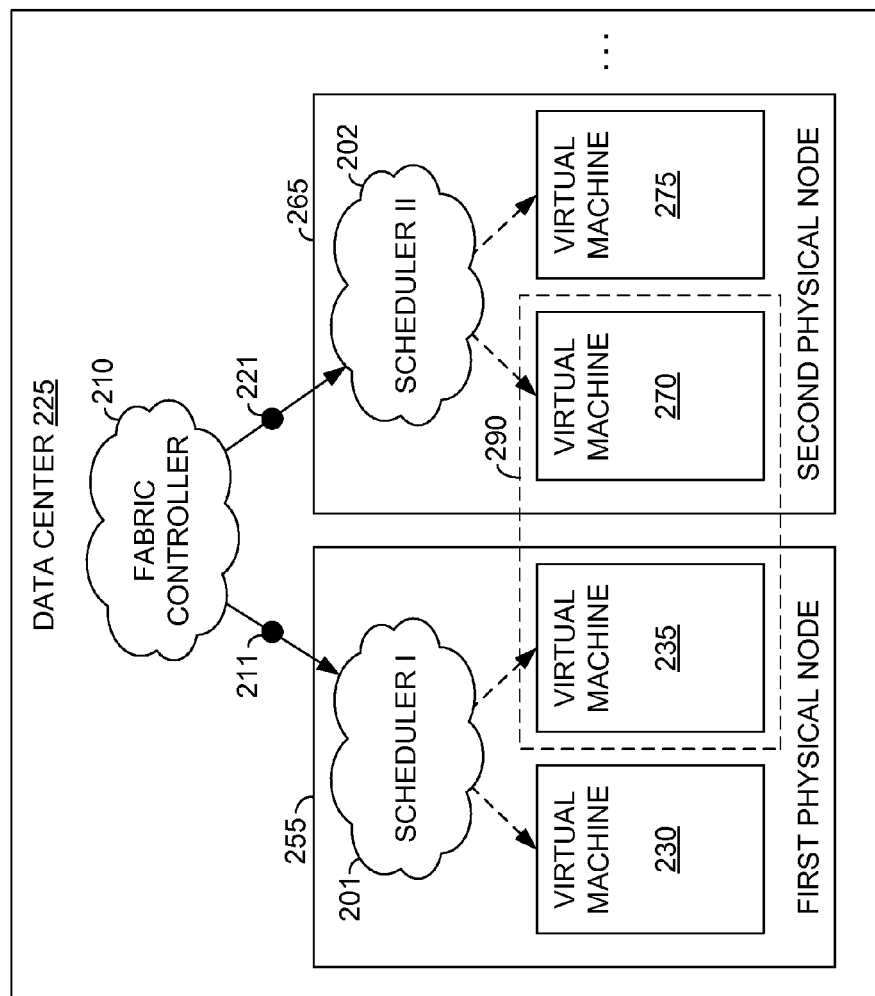


FIG. 2.

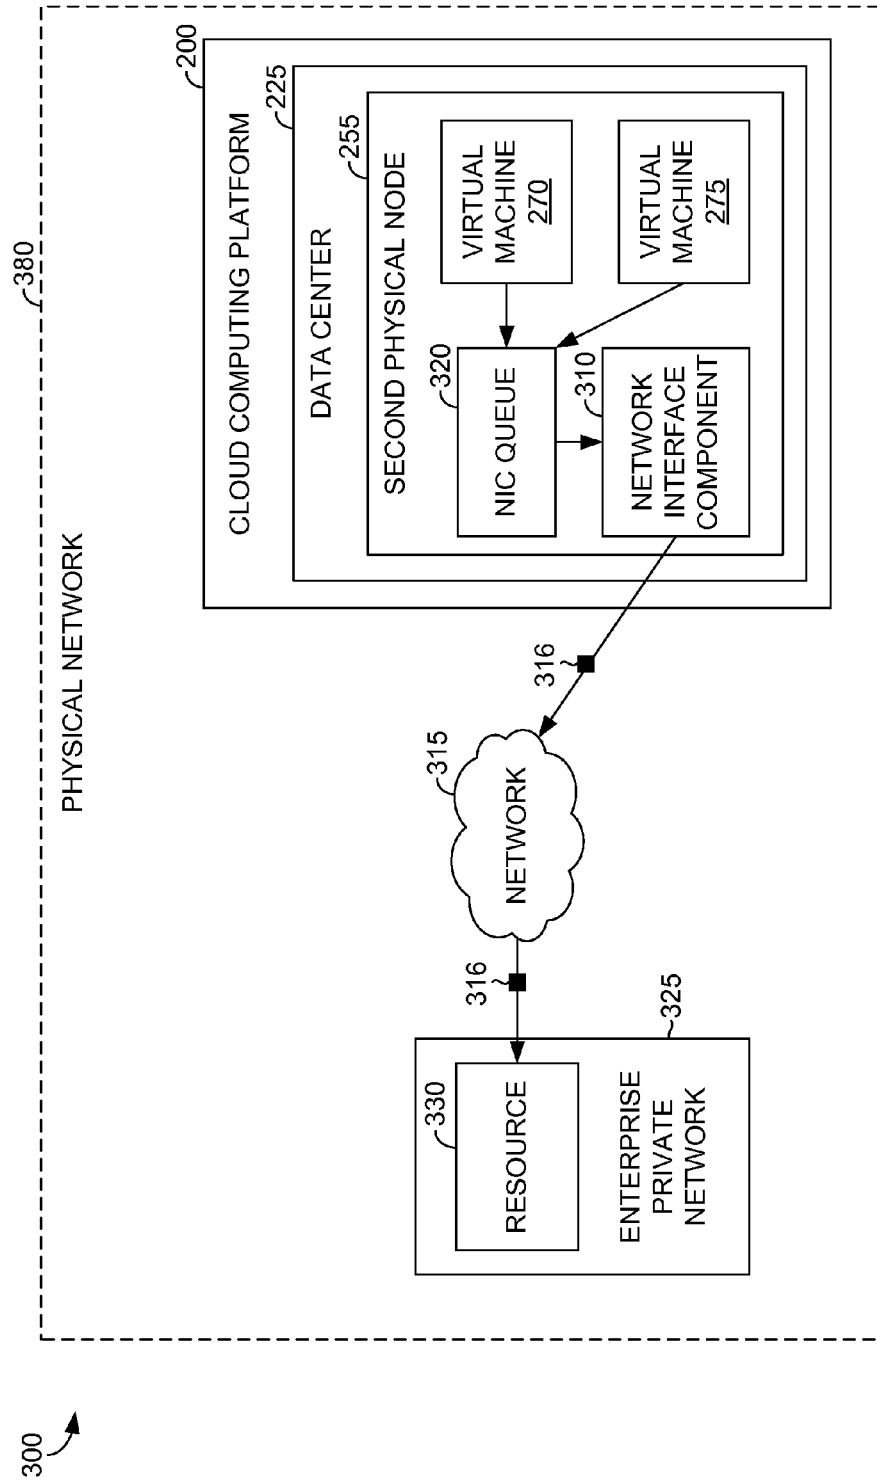


FIG. 3.

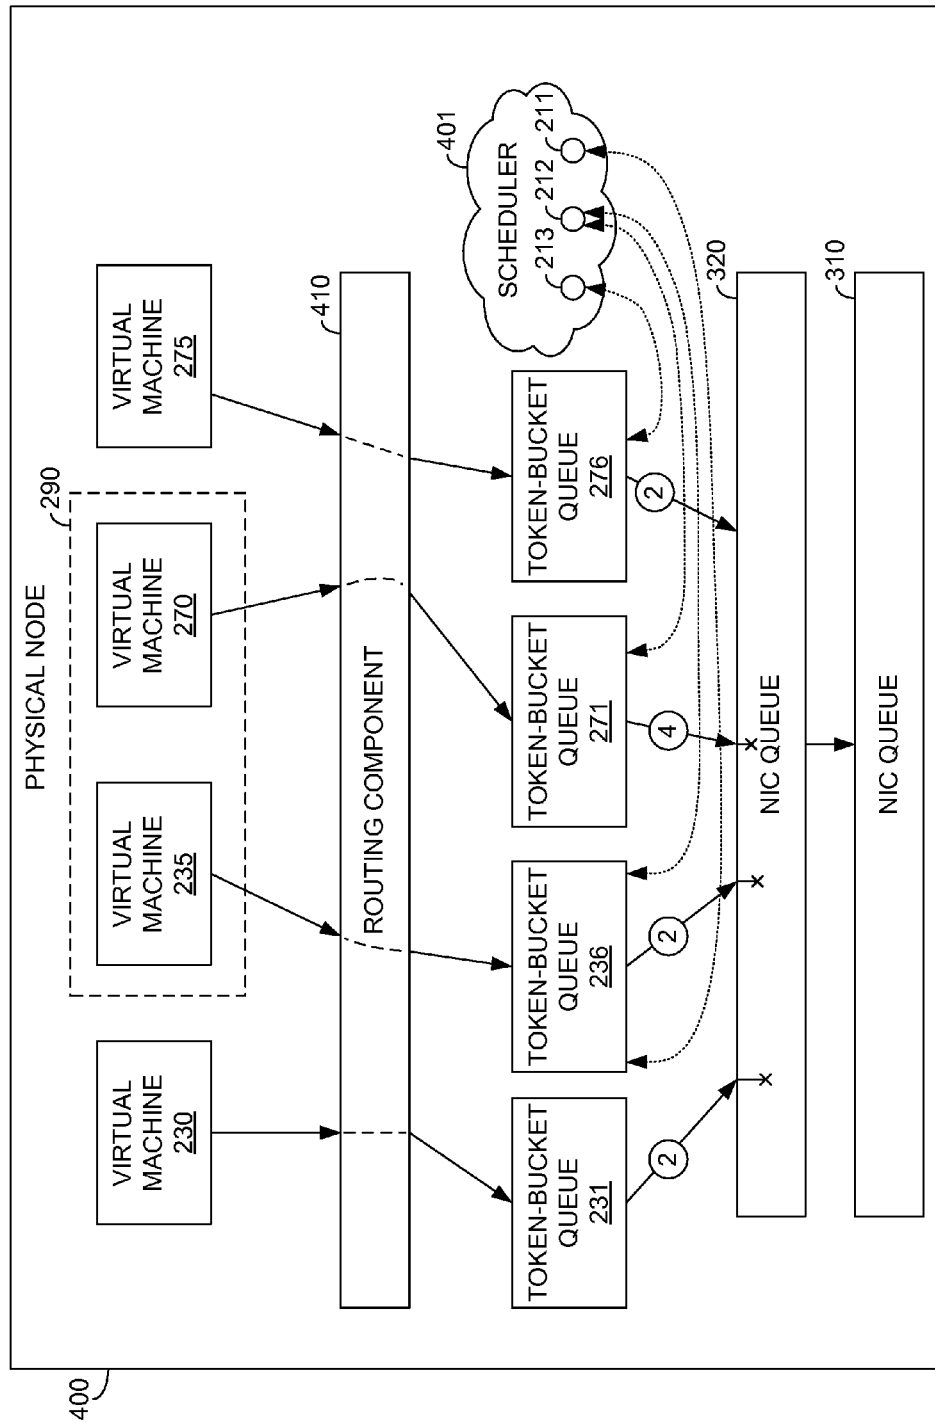


FIG. 4.

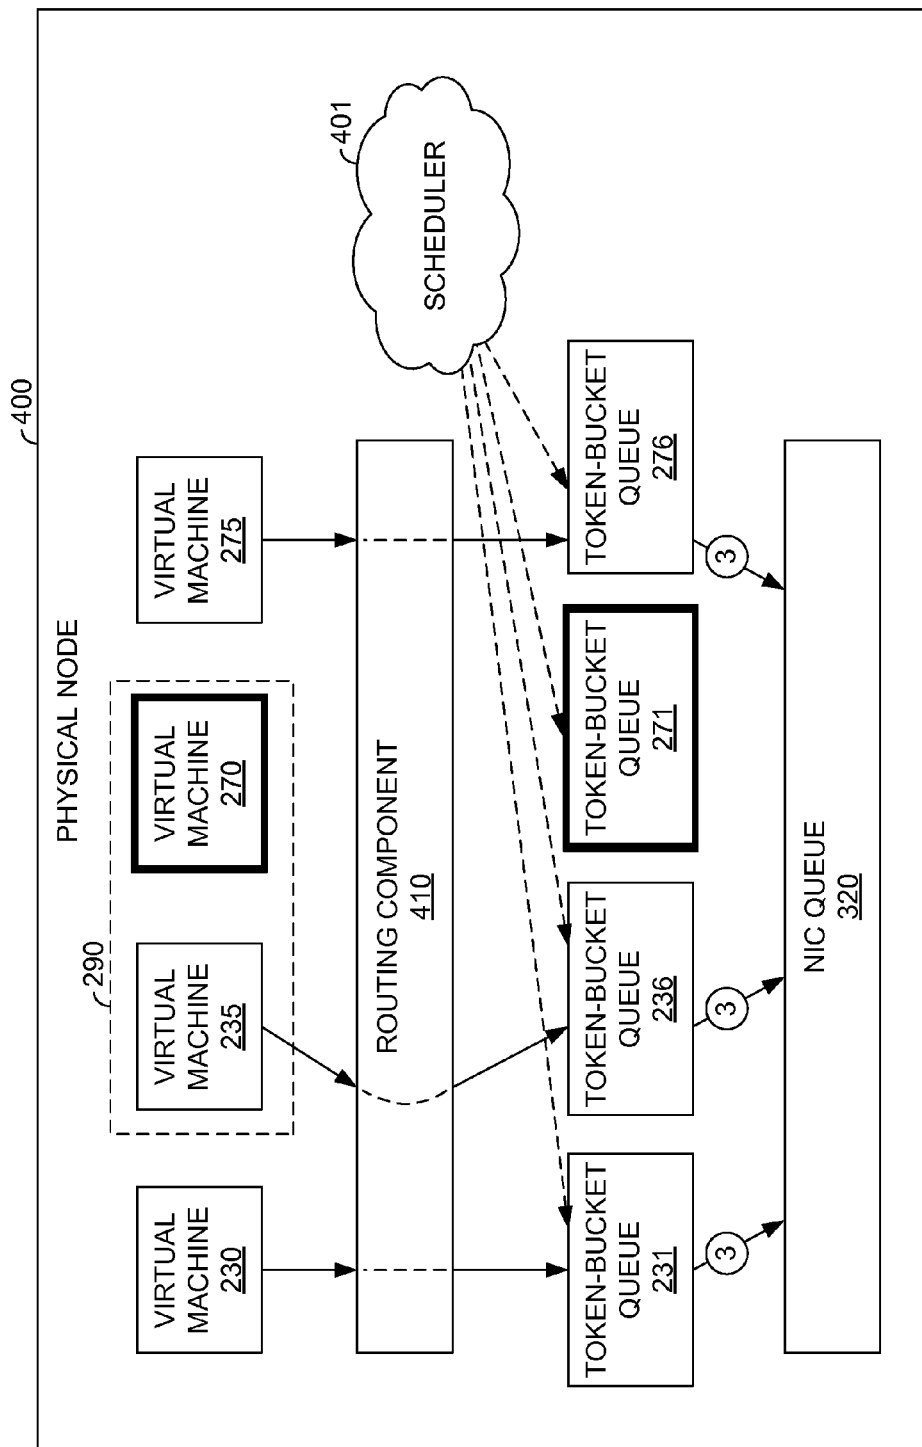


FIG. 5.

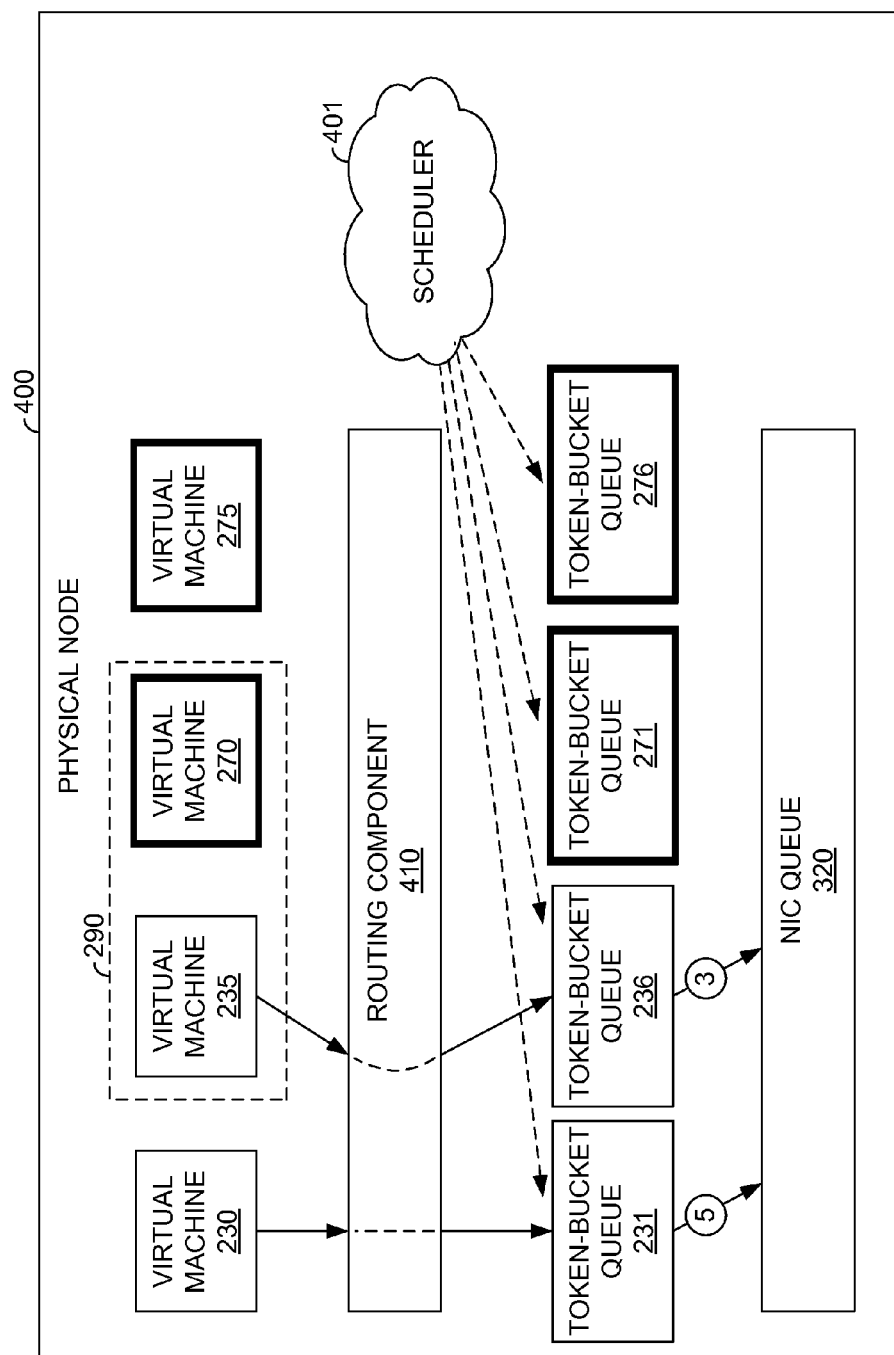


FIG. 6.

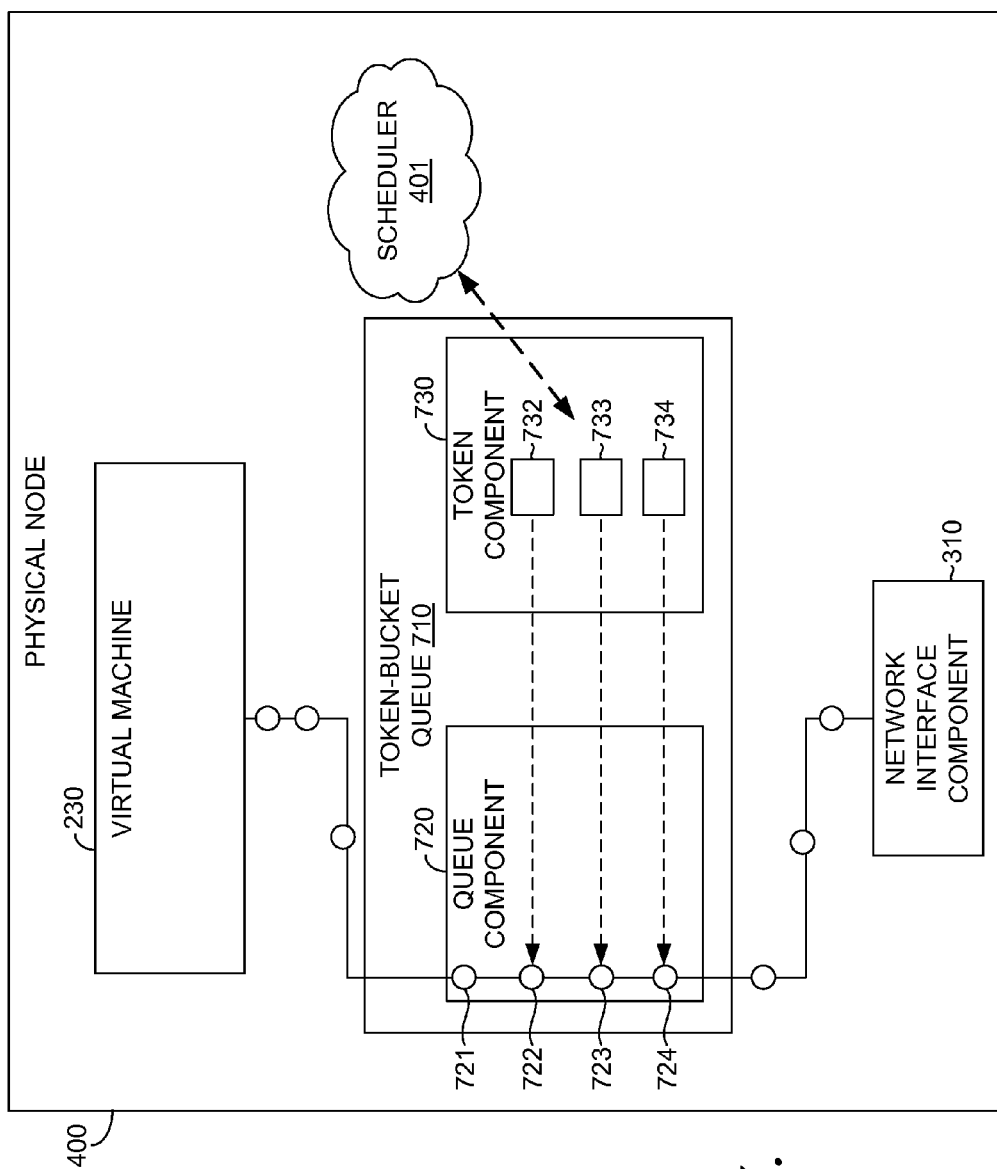


FIG. 7.

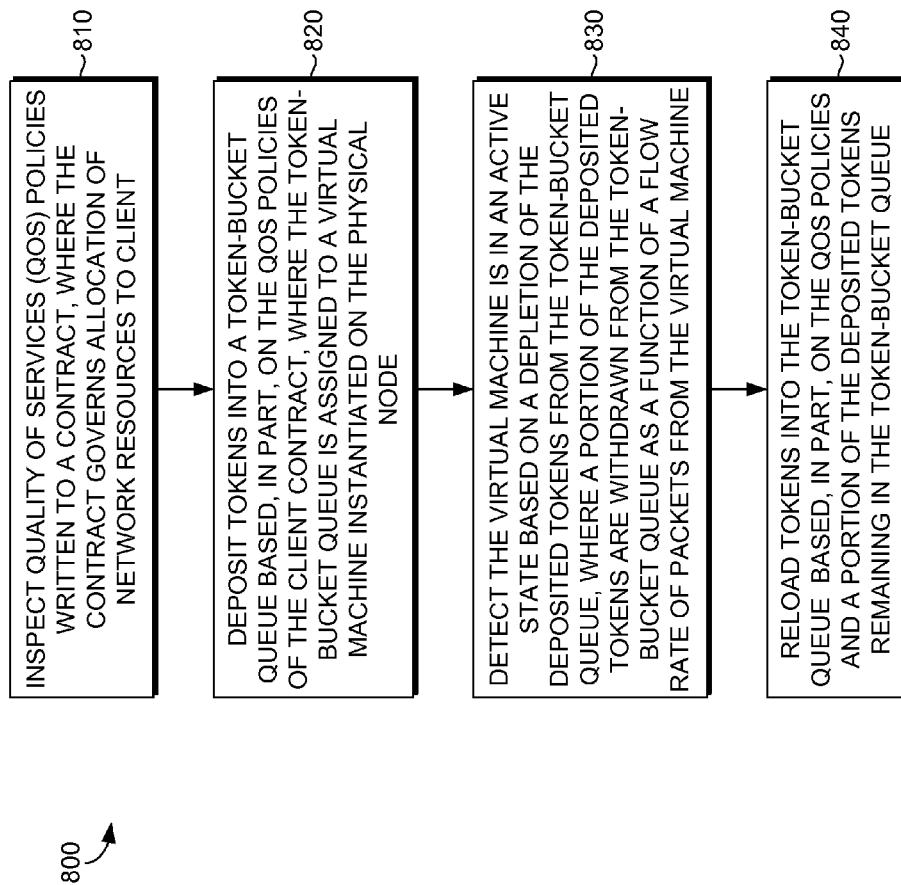


FIG. 8.

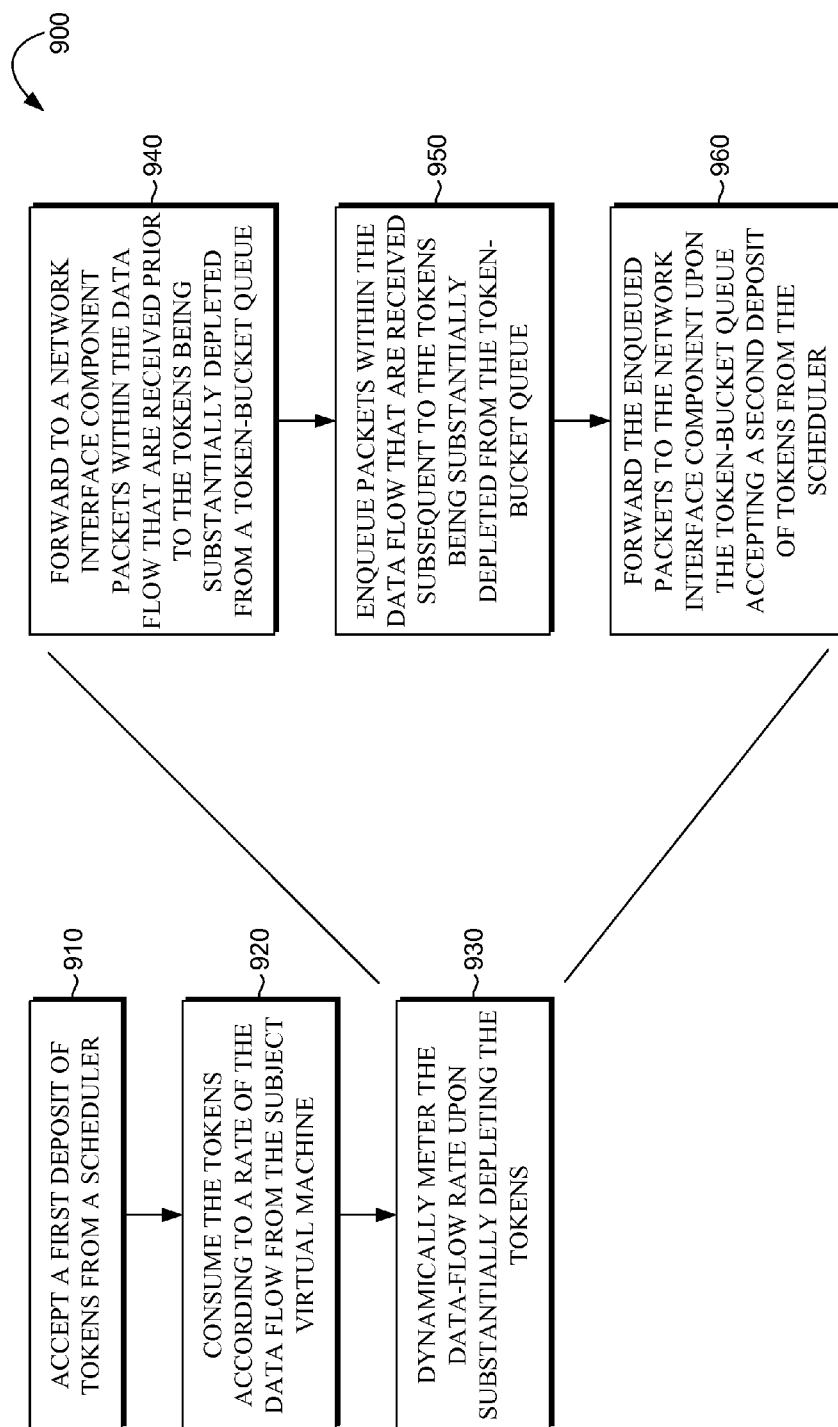


FIG. 9.

1

# APPLYING POLICIES TO SCHEDULE NETWORK BANDWIDTH AMONG VIRTUAL MACHINES

## CROSS REFERENCE TO RELATED APPLICATIONS

This application is a continuation of, and claims priority from, U.S. patent application Ser. No. 12/790,981, filed May 31, 2010, entitled "APPLYING POLICIES TO SCHEDULE NETWORK BANDWIDTH AMONG VIRTUAL MACHINES," which is incorporated herein by reference in its entirety.

## BACKGROUND

Large-scale networked systems are commonplace platforms employed in a variety of settings for running applications and maintaining data for business and operational functions. For instance, a data center (e.g., physical cloud computing infrastructure) may provide a variety of services (e.g., web applications, email services, search engine services, etc.) for a plurality of customers simultaneously. These large-scale networked systems typically include a large number of resources distributed throughout the data center, in which each resource resembles a physical machine or a virtual machine (VM) running on a physical node or host. When the data center hosts multiple tenants (e.g., customer programs), these resources are optimally allocated from the same data center to the different tenants.

Often, multiple VMs will concurrently run on the same physical node within a computing network, or the data center. These VMs that share a common physical node may be allocated to the different tenants and may require different amounts of resources at a various times. For instance, there is typically only one network card installed to a physical node, where the network card has a limited amount of network bandwidth. When one tenant requires a high amount of resources to accomplish a particular task, one of the tenant's VMs running on the physical node can potentially stress the network card by sending a large amount of packets thereover, preventing other VMs running on the physical node from fairly sharing the resources thereof.

Accordingly, because optimal allocation of processing capabilities within a data center often involves placing a plurality of VMs on a single physical node, and because these VMs require resources (e.g., network bandwidth) to carry out most operations, a mechanism that allows the VMs on the same physical node to share the network bandwidth fairly, by placing restrictions on resource consumption, would help reduce stressing available resources and would help maintain a high quality of service provided to the tenants of the data center.

## SUMMARY

This Summary is provided to introduce concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

Embodiments of the present invention provide a mechanism that allows the virtual machines (VMs) residing on a common physical node to fairly share the resources (e.g., network bandwidth) provided by the physical node. Typically, fairly sharing resources involves placing restrictions

2

on resource consumption to ameliorate stressing the network bandwidth or adversely affecting the quality of service that is guaranteed to tenants of the data center in which the VMs reside.

In an exemplary embodiment, the placement of restrictions may involve providing a management utility (e.g., scheduler **401** of FIGS. **5-7**) that dynamically controls VM experiences by allocating the networking bandwidth to each of the VMs on a shared physical machine as a function of a policy. In other words, the management utility may grant a physical machine owner, or a tenant of the data center, the ability to configure the policies that control the VMs usage of network bandwidth. Beyond serving as an interface that accepts policies being dictated by one or more entities, the management utility may enforce these policies so that the VMs are provided with a certain quality of service (e.g., a minimum/maximum network bandwidth guaranteed for each VM).

In one instance of enforcing these policies, the management utility may employ rate limiting techniques that are used to control a volume of traffic being sent from the VMs. By way of example, one rate limiting technique may instruct that data packets be freely conveyed from a VM to a network when a rate of data flow from the VM is consuming less resources than currently allocated to the VM. Generally, the allocated resources are specified by the policies associated with the VM. Another rate limiting technique may instruct that data packets be dropped or delayed when the VM is attempting to utilize more resources than currently allocated thereto.

As more fully described below, a set of token-bucket queues and a scheduler may be implemented on the physical node to carry out these rate limiting techniques. Carrying out the rate limiting techniques may initially involve the scheduler inspecting one or more quality of service (QoS) policies written to a contract, where the contract governs allocation of network resources to a client or tenant of a data center. Next, the scheduler may deposit tokens into a token-bucket queue based, in part, on the QoS policies of the client contract. As described herein, the token-bucket queue is assigned to a VM instantiated on the physical node that is shared by other VMs.

The scheduler may then inspect the token-bucket queue to ascertain a state of the VM. For instance, when the deposited tokens are not depleted from the token-bucket queue, the schedule may detect that the VM is in an idle state. In contrast, the scheduler may detect that the VM is in an active state based on a depletion of the deposited tokens from the token-bucket queue, as portion of the deposited tokens are withdrawn from the token-bucket queue as a function of a rate of the data flow from the VM. In this instance, the scheduler may reload tokens into the token-bucket queue based on the QoS policies and on a portion of the deposited tokens remaining in the token-bucket queue. As such, the allocation of network bandwidth using the scheduler in conjunction with the token-bucket queues maintains fairness amongst the VMs while maximizing throughput on the physical node.

## BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention are described in detail below with reference to the attached drawing figures, wherein:

FIG. **1** is a block diagram of an exemplary computing environment suitable for use in implementing embodiments of the present invention;

3

FIG. 2 is a block diagram illustrating an exemplary cloud computing platform, suitable for use in implementing embodiments of the present invention, that is configured to allocate bandwidth within a physical node of a data center;

FIG. 3 is block diagram of an exemplary distributed computing environment depicting components within a physical node communicating to an external endpoint, in accordance with an embodiment of the present invention;

FIGS. 4-6 are block diagrams depicting various data flows between virtual machines (VMs) and corresponding token-bucket queues residing on a single physical node, in accordance with an embodiment of the present invention;

FIG. 7 is a block diagram of a token-bucket queue that is metering a data flow based on tokens deposited therein, in accordance with embodiments of the present invention;

FIG. 8 is a flow diagram showing an exemplary method for metering various data flows of packets originating from a plurality of VMs instantiated on the single physical node, in accordance with an embodiment of the present invention; and

FIG. 9 is a flow diagram showing an exemplary method for imposing a limit on a rate of data flowing from a subject VM, in accordance with an embodiment of the present invention.

#### DETAILED DESCRIPTION

The subject matter of embodiments of the present invention is described with specificity herein to meet statutory requirements. However, the description itself is not intended to limit the scope of this patent. Rather, the inventors have contemplated that the claimed subject matter might also be embodied in other ways, to include different steps or combinations of steps similar to the ones described in this document, in conjunction with other present or future technologies. Moreover, although the terms “step” and/or “block” may be used herein to connote different elements of methods employed, the terms should not be interpreted as implying any particular order among or between various steps herein disclosed unless and except when the order of individual steps is explicitly described.

Embodiments of the present invention relate to methods, computer systems, and computer-readable media for applying a throttle to a source (e.g., virtual machines (VMs) instantiated on a physical node) of packets, as opposed to controlling traffic at network edges. The applied throttle may, in operation, meter data flows from the VMs according to contracts associated with each. In an exemplary embodiment, metering involves delaying delivery of those packets that exceed one or more policies articulated by the contracts. By way of example, metering may be implemented by token-bucket queues that temporarily store, or “enqueue,” the delayed packets. Typically, the token-bucket queues act as first-in first-out (FIFO) buffers that hold the delayed packets until network bandwidth is allocated to the VMs that originated the delayed packets. Upon allocation of sufficient network bandwidth, the token-bucket queues are designed to release the delayed packets to be delivered to their intended destination.

In one aspect, embodiments of the present invention relate to one or more computer-readable media having computer-executable instructions embodied thereon that, when executed, perform a method for metering various data flows of packets originating from a plurality of VMs instantiated on the single physical node. Initially, the method includes the step of inspecting one or more quality of service (QoS) policies written to a contract. Typically, the contract governs

4

allocation of network resources to a client. Next, tokens may be deposited into a token-bucket queue based, in part, on the one or more QoS policies of the client contract. As shown in FIGS. 4-6, the token-bucket queue is assigned to the VM instantiated on the physical node. The VM may be identified as having assumed an active state based on a depletion of the deposited tokens from the token-bucket queue. (When in the active state, a portion of the deposited tokens are withdrawn from the token-bucket queue as a function of a rate of the data flow from the VM.) The method may further include the step of reloading tokens into the token-bucket queue based, in part, on (a) the QoS policies and/or (b) a portion of the deposited tokens remaining in the token-bucket queue.

In another aspect, embodiments of the present invention relate to a computer system capable of metering a rate at which packets from a VM are injected into a network utilizing one or more policies specified by tenants of a data center. The computer system includes a computer storage medium that has computer software components embodied thereon. The computer software components include the following: a first VM, a second VM, a first token-bucket queue, a second token-bucket queue, and a scheduler. The first and second VMs are instantiated on the same physical node. The first token-bucket queue is assigned to the first VM, while the second token-bucket queue is assigned to the second VM. As used herein, the first token-bucket queue and the second token-bucket queue generally represent distinct memory buffers located on the physical node.

The scheduler is configured to read a first-client contract and a second-client contract associated with the first VM and the second VM, respectively. The scheduler is further able to deposit a first amount of tokens into the first token-bucket queue as a function of policies written to the first-client contract. Similarly, the scheduler may deposit a second amount of tokens into the second token-bucket queue as a function of the policies within the second-client contract. As used herein, the first amount of tokens and the second amount of tokens each generally represent a portion of network bandwidth that is allocated to the first VM and the second VM, respectively. Often, the first amount and the second amount of tokens differ in value based on a discrepancy in network bandwidth allocation between the first and second VM.

In yet another aspect, embodiments of the present invention relate to a computerized method for imposing a limit on a rate of data flowing from a subject VM. In one embodiment, the method involves accepting a first deposit of tokens from a scheduler. Typically, an amount of tokens provided within the first deposit is based on, at least one of, the following: tokens currently maintained within a token-bucket queue, network bandwidth allocated to the subject VM associated with the token-bucket queue, or network bandwidth allocated to other VMs located on a common physical node as the subject VM.

The method may further involve the steps of consuming the tokens according to a rate of the data flow from the subject VM, and dynamically metering the data-flow rate upon substantially depleting the tokens. In instances of the present invention, the process of metering may include either forwarding to a network interface component packets within the data flow that are received prior to the tokens being substantially depleted from the token-bucket queue, or enqueueing packets within the data flow that are received subsequent to the tokens being substantially depleted from the token-bucket queue. Upon the token-bucket queue accepting a second deposit of tokens from the scheduler, the enqueued packets may be forwarded to the network interface

5

component, assuming that the tokens within the second deposit is adequate in amount to distribute the packets held by the token-bucket queue.

Having briefly described an overview of embodiments of the present invention, an exemplary operating environment suitable for implementing embodiments of the present invention is described below. In particular, an exemplary operating environment supports functionality of a mechanism that allows VMs residing on a common physical node to fairly share network bandwidth by placing restrictions on resource consumption; thus, ameliorating overly depleting the network bandwidth or adversely affecting the quality of service that is guaranteed to the VMs via service contracts.

Referring to the drawings in general, and initially to FIG. 1 in particular, an exemplary operating environment for implementing embodiments of the present invention is shown and designated generally as computing device 100. Computing device 100 is but one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of embodiments of the present invention. Neither should the computing environment 100 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated.

Embodiments of the present invention may be described in the general context of computer code or machine-useable instructions, including computer-executable instructions such as program components, being executed by a computer or other machine, such as a personal data assistant or other handheld device. Generally, program components including routines, programs, objects, components, data structures, and the like refer to code that performs particular tasks, or implements particular abstract data types. Embodiments of the present invention may be practiced in a variety of system configurations, including handheld devices, consumer electronics, general-purpose computers, specialty computing devices, etc. Embodiments of the invention may also be practiced in distributed computing environments where tasks are performed by remote-processing devices that are linked through a communications network.

With continued reference to FIG. 1, computing device 100 includes a bus 110 that directly or indirectly couples the following devices: memory 112, one or more processors 114, one or more presentation components 116, input/output (I/O) ports 118, I/O components 120, and an illustrative power supply 122. Bus 110 represents what may be one or more busses (such as an address bus, data bus, or combination thereof). Although the various blocks of FIG. 1 are shown with lines for the sake of clarity, in reality, delineating various components is not so clear, and metaphorically, the lines would more accurately be grey and fuzzy. For example, one may consider a presentation component such as a display device to be an I/O component. Also, processors have memory. The inventors hereof recognize that such is the nature of the art and reiterate that the diagram of FIG. 1 is merely illustrative of an exemplary computing device that can be used in connection with one or more embodiments of the present invention. Distinction is not made between such categories as “workstation,” “server,” “laptop,” “handheld device,” etc., as all are contemplated within the scope of FIG. 1 and reference to “computer” or “computing device.”

Computing device 100 typically includes a variety of computer-readable media. By way of example, and not limitation, computer-readable media may comprise Random Access Memory (RAM); Read Only Memory (ROM); Electronically Erasable Programmable Read Only Memory (EEPROM); flash memory or other memory technologies;

6

CDROM, digital versatile disks (DVDs) or other optical or holographic media; magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium that can be used to encode desired information and be accessed by computing device 100.

Memory 112 includes computer storage media in the form of volatile and/or nonvolatile memory. The memory may be removable, nonremovable, or a combination thereof. Exemplary hardware devices include solid-state memory, hard drives, optical-disc drives, etc. Computing device 100 includes one or more processors that read data from various entities such as memory 112 or I/O components 120. Presentation component(s) 116 present data indications to a user or other device. Exemplary presentation components include a display device, speaker, printing component, vibrating component, etc. I/O ports 118 allow computing device 100 to be logically coupled to other devices including I/O components 120, some of which may be built-in. Illustrative components include a microphone, joystick, game pad, satellite dish, scanner, printer, wireless device, etc.

With reference to FIGS. 1 and 2, a first physical node 255 and/or second physical node 265 may be implemented by the exemplary computing device 100 of FIG. 1. Further, one or more of the virtual machines 230, 235, 270, and/or 275 may include portions of the memory 112 of FIG. 1 and/or portions of the processors 114 of FIG. 1. Similarly, the token-bucket queues 231, 236, 271, and/or 276 of FIGS. 4-6 may include portions of the memory 112 of FIG. 1 and/or portions of the processors 114 of FIG. 1.

Turning now to FIG. 2, a block diagram is illustrated, in accordance with an embodiment of the present invention, showing an exemplary cloud computing platform 200 that is configured to allocate network bandwidth to the virtual machines 230, 235, 270 and 275 within a data center 225 for use by various service applications. By way of example, service application 290 includes virtual machines 235 and 270 that are instantiated on physical nodes 255 and 256, respectively, to support functionality of the service application 290. It will be understood and appreciated that the cloud computing platform 200 shown in FIG. 2 is merely an example of one suitable computing system environment and is not intended to suggest any limitation as to the scope of use or functionality of embodiments of the present invention. For instance, the cloud computing platform 200 may be a public cloud, a private cloud, or a dedicated cloud. Neither should the cloud computing platform 200 be interpreted as having any dependency or requirement related to any single component or combination of components illustrated therein. Further, although the various blocks of FIG. 2 are shown with lines for the sake of clarity, in reality, delineating various components is not so clear, and metaphorically, the lines would more accurately be grey and fuzzy. In addition, any number of physical machines, virtual machines, data centers, endpoints, or combinations thereof may be employed to achieve the desired functionality within the scope of embodiments of the present invention.

The cloud computing platform 200 includes the data center 225 configured to host and support operation of the virtual machines 230, 235, 270, and 275, or endpoints, of service applications (e.g., service application 290) owned by tenants/customers of the data center 225. The phrase “service application,” as used herein, broadly refers to any software, or portions of software, that runs on top of, or accesses storage locations within, the data center 225. In one embodiment, one or more of the virtual machines 230, 235, 270, and 275 may represent the portions of software, component programs, or instances of roles that participate in the

service application. In another embodiment, one or more of the virtual machines **230**, **235**, **270**, and **275** may represent stored data that is accessible to a service application. It will be understood and appreciated that the virtual machines **230**, **235**, **270**, and **275** shown in FIG. 2 are merely an example of suitable parts to support a service application and are not intended to suggest any limitation as to the scope of use or functionality of embodiments of the present invention.

Generally, virtual machines **230**, **235**, **270**, and **275** are instantiated and designated to a service application based on demands (e.g., amount of processing load) placed on the service application. As used herein, the phrase “virtual machine” is not meant to be limiting, and may refer to any software, application, operating system, or program that is executed by a processing unit to underlie the functionality of the service application. Further, the virtual machines **230**, **235**, **270**, and **275** may include processing capacity, storage locations, and other assets within the data center **225** to properly support the service applications.

In operation, the virtual machines **230**, **235**, **270**, and **275** are dynamically allocated resources (e.g., network bandwidth of a network interface card) of the physical nodes **255** and **265** to satisfy a current processing load while meeting QoS policies **211** and **221** of client contracts. In one instance, a fabric controller **210** is responsible for automatically pushing down these QoS policies **211** and **221** to schedulers **201** and **202**, respectively. By way of example, the fabric controller **210** may rely on a service model (e.g., designed by a customer that owns the service application) to provide guidance on how and when to allocate resources, such as processing capacity, to the virtual machines **230**, **235**, **270**, and **275**. In another example, the fabric controller **210** may be configured to parse service-level agreements (SLAs) to identify which guarantees are provided to which customers and to which components running the customer’s service application. These guarantees may be delivered to the schedulers **201** and **202** in the form of the QoS policies **211** and **221**. In operation, the schedulers **201** and **202** enforce the QoS policies **211** and **221** via the token-bucket queues (discussed below).

As used herein, the term “scheduler” is not meant to be limiting, but may encompass any mechanism that promotes network-bandwidth management by receiving instructions from a centralized utility (e.g., fabric controller **210**) and enforcing QoS policies within the instructions. In an exemplary embodiment, the schedulers **201** and **202** may monitor data flows originating from each of the VMs on the physical nodes **255** and **265**, and compare the data flows against the QoS policies to arrive upon an appropriate amount of network bandwidth to allocate to the virtual machines **230**, **235**, **270**, and **275**. In the field of computer networking (e.g., packet-switched telecommunication networks), the traffic engineering phrase “QoS policies” refers to a blueprint for implementing resource reservation and distribution amongst VMs. Generally, “quality of service” (QoS) refers to the ability to provide different priority to different service applications, components of applications (e.g., VMs), users, or data flows, in order to guarantee a certain level of performance thereto. For example, a fixed bit rate may be guaranteed for a particular VM in a particular context. As such, the guarantees of the QoS policies **211** and **221** are advantageously utilized as guidelines for scheduling resources to users, who share the same network resources within a limited-capacity physical node.

As discussed above, one or more VMs may be dynamically allocated resources within the first physical node **255** and the second physical node **265**. Per embodiments of the

present invention, the physical nodes **255** and **265** may represent any form of computing devices, such as, for example, a personal computer, a desktop computer, a laptop computer, a mobile device, a consumer electronic device, server(s), the computing device **100** of FIG. 1, and the like. In one instance, the physical nodes **255** and **265** host and support the operations of the virtual machines **230**, **235**, **270**, and **275**, respectively, while simultaneously hosting other virtual machines (not shown) carved out for supporting other tenants of the data center **225**, where the tenants include endpoints of other service applications owned by different customers.

In one aspect, the VMs operate within the context of the cloud computing platform **200** and, accordingly, may communicate internally within a physical node, communicate across physical nodes **255** and **265** through connections dynamically made that traverse the data center **225**, and externally through a physical network topology to resources of a remote network (e.g., in FIG. 3 resource **375** of the enterprise private network **325**). The connections may involve interconnecting the VMs distributed across physical resources of the data center **225** via a network cloud (not shown). The network cloud interconnects these resources such that one VM may recognize a location of another VM in order to establish a communication therebetween. For instance, the network cloud may establish this communication over channels connecting the VMs **235** and **270** of the service application **290**. By way of example, the channels may include, without limitation, one or more local area networks (LANs) and/or wide area networks (WANs). Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets, and the Internet. Accordingly, the network is not further described herein.

Turning now to FIG. 3, block diagram illustrating an exemplary distributed computing environment **300** with two VMs **270** and **275** established therein is shown in accordance with an embodiment of the present invention. Initially, the distributed computing environment **300** includes a physical network **380** that includes an enterprise private network **325** and the cloud computing platform **200**, as discussed with reference to FIG. 2. As used herein, the phrase “physical network” is not meant to be limiting, but may encompass tangible mechanisms and equipment (e.g., fiber lines, circuit boxes, switches, antennas, IP routers, and the like), as well as intangible communications and carrier waves, that facilitate communication between endpoints at geographically remote locations. By way of example, the physical network **380** may include any wired or wireless technology utilized within the Internet, or available for promoting communication between disparate networks.

Generally, customers of the data center **225** often require service applications running on resources therein to interact with the software being run at another location, such as the private enterprise network **325**. As used herein, the phrase “private enterprise network” generally refers to a server managed by a customer, or client of the cloud computing platform **200**, that is geographically remote from the data center **225**. In the embodiment illustrated in FIG. 3, the enterprise private network **325** includes resources, such as resource **330**, that are managed by the client of the cloud computing platform **200**. Often, these resources host and support operations of components of the service application owned by the customer.

As depicted by reference numeral **315**, a network is established within the physical network **380**. Typically, the network **315** is established to transport packets between endpoints. By way of example, the network **315** acts to

convey packets 316 between the VMs 270 and 275 instantiated on the second physical node 255 and the resource 330. In operation, both the VMs 270 and 275 may originate data flows that include packets. These data flows may be intercepted by a network interface controller (NIC) queue that meters a first data flow from the virtual machine 270 in conjunction with a second data flow from the virtual machine 275. Generally, the NIC queue 320 functions as an internal filter, which controllably forwards the packets of the first and second data flows to a network interface component 310.

Upon passing the metered data flow to the network interface component 310, which may be configured as a network adapter or network interface card, the network interface component 310 releases the packets 316 of the data flows into the network 315. Stated differently, the network interface component 310 receives the metered first data flow and the metered second data flow from the NIC queue 320 and releases from the second physical node 255 a controlled volume of traffic comprising the packets 316 within the metered first and second data flows. In other embodiments (not shown), virtualized implementations of the network interface controller (VNICs) may be set up within the virtual machines 270 and 275 themselves to self-regulate a rate of outgoing packets.

As mentioned above, the network 315 is limited by its available network bandwidth while the second physical node 255 is configured with a particular amount of resources (e.g., limited processing capacity of the network interface component 310). Accordingly, if no additional mechanisms to fairly distribute the network bandwidth and/or resources are put in place, the VMs 270 and 275 may compete for the network bandwidth and/or resources without consideration of predefined QoS policies. For instance, once a connection over the network 315 is established, the virtual machine 270 may attempt consume all the network bandwidth independent of other usage, thus, degrading performance of the virtual machine 275, which is also attempting to send packet(s) to an endpoint. As such, exemplary embodiments of the present invention introduce token-bucket queues assigned to each VM that meter a rate at which packets from the VMs are injected into the network 315 utilizing the QoS policies specified by tenants of the data center 225. These token-bucket queues are described immediately below with reference to FIGS. 4-7.

Typically, the resource 330 and the data center 225 include, or are linked to, some form of a computing unit (e.g., central processing unit, microprocessor, etc.) to support operations of the VM(s) 270 and 275, and/or component(s), running thereon. As utilized herein, the phrase "computing unit" generally refers to a dedicated computing device with processing power and storage memory, which supports one or more operating systems or other underlying software. In one instance, the computing unit is configured with tangible hardware elements, or machines, that are integral, or operably coupled, to the resource 330 and the data center 225 to enable each device to perform a variety of processes and operations. In another instance, the computing unit may encompass a processor (not shown) coupled to the computer-readable medium accommodated by each of the resource 330 and the data center 225. Generally, the computer-readable medium stores, at least temporarily, a plurality of computer software components (e.g., the VMs 270 and 275) that are executable by the processor. As utilized herein, the term "processor" is not meant to be limiting and may encompass any elements of the computing unit that act in a computational capacity.

In such capacity, the processor may be configured as a tangible article that processes instructions. In an exemplary embodiment, processing may involve fetching, decoding/interpreting, executing, and writing back instructions.

Turning now to FIG. 4, a block diagram is shown depicting various data flows between VMs 230, 235, 270, and 275 and corresponding token-bucket queues 231, 236, 271, and 276, respectively, in accordance with an embodiment of the present invention. Initially, this plurality of VMs each reside on a single physical node 400. Accordingly, the VMs 230, 235, 270, and 275 are obliged to fairly share network resources, such as a network link/connection or the bandwidth provisioned thereto.

A routing component 410 is provided that knows a mapping between the VMs 230, 235, 270, and 275 and their assigned token-bucket queues 231, 236, 271, and 276, respectively. In operation the routing component 410 is configured to access the mapping upon detecting one or more VMs attempting to transmit a data flow to an endpoint external to the physical node 400, and to route the data flow to the one or more token-bucket queues based on the mapping. By way of example, the routing component 410 may capture a data flow, identify a source of the data flow as the virtual machine 230, and utilize the mapping to determine that the data flow should be passed to the token-bucket queue 231, which is assigned to the virtual machine 230.

The above described embodiment typically applies when a packets within the data flow are addressed to a destination outside the physical node 400. If the packets within the data flow are addressed to a destination internal to the physical node 400, the routing component 410 may directly route the packets to the targeted VM. By way of example, if the VM 230 is attempting to communicate packets to the VM 235, the routing component 410 may convey the packets to the VM 235 without passing through the token-bucket queue 231 assigned to the VM 230, thereby bypassing and precluding engagement of the token-bucket queue 231. This is due to the lack of network-bandwidth consumption when internally transmitting packets.

An alternative design is to locate the token-bucket queues 231, 236, 271, and 276 above the routing component 410. As such, the VM traffic will pass through one or more of the token-bucket queues 231, 236, 271, and 276 first and then the routing component 410. In that way, the token-bucket queues can control the QoS of the traffic between the VMs 230, 235, 270, and 275 and in addition to the traffic that is transmitted to an external client. But, the traffic targeting on internal clients and on external clients will be treated similarly and controlled by the same policy, as the destination is not known when the packets pass through the token-bucket queues 231, 236, 271, and 276. Or, even another token-bucket queue (not shown) can be added at the inbound path of each of the VMs 230, 235, 270, and 275. Along with the original token-bucket queues token-bucket queues 231, 236, 271, and 276 above the NIC queue 310, they can control the internal and external traffic of each VM separately by applying different policies.

A scheduler 401 is provided on the physical node 400. Generally, the scheduler 401 represents local VM management software that dynamically allocates network bandwidth to the VMs 230, 235, 270, and 275 based on a variety of criteria. In embodiments, schedulers are instantiated on physical nodes by way of a one-to-one relationship. In embodiments, each node in a data center is provided with a separate scheduler and these separate schedulers will work independently. By way of example, two schedulers residing

on two separate nodes will not coordinate with each other. As mentioned above, the scheduler 401 maintains QoS policies 211, 212, and 213 that are pushed down from a fabric controller (not shown) that maintains prevailing client, or subscriber, contracts containing the QoS policies 211, 212, and 213. In addition to the QoS policies 211, 212, and 213, other information (e.g., changes to QoS policies) may migrate from the fabric controller to the scheduler 401 at various periods of time (e.g., regular intervals or upon receiving an update). In one instance, updates to the client contracts may be made via a user-mode API that allows an administrator to specify the QoS policies for the VMs 230, 235, 270, and 275. In another example, a value of relative weighting or absolute bandwidth may be entered at the user-mode API. The “relative weighting” and the “absolute bandwidth” are variables used in independent, mutually-exclusive schemes for allocating network bandwidth, which are discussed in more detail below.

In operation, the scheduler 401 is capable of inspecting the QoS policies 211, 212, and 213 written to one or more clients’ contracts. As discussed above, the client contracts govern allocation of network resources to the VMs 230, 235, 270, and 275. (Because VM 235 and VM 270 are components of a common service application (see reference numeral 290 of FIG. 2), the QoS policies 212 govern distribution of network bandwidth to both the VMs 235 and 270.) Upon inspection of the QoS policies 211, 212, and 213, the scheduler 401 may deposit tokens into the token-bucket queues 231, 236, 271, and 276 based, in part, on the QoS policies 211, 212, and 213. As used herein, the term “tokens” is not meant to be restrictive, and can encompass any representations of a portion of network bandwidth that is allocated to one or more VMs.

The token-bucket queues 231, 236, 271, and 276 are provided to impose a limit on a rate of data flowing from their corresponding VMs 230, 235, 270, and 275, respectively, as a function of the deposited tokens. The process of imposing a limit of the data-flow rate will now be discussed with reference to FIGS. 7 and 9. FIG. 9 shows a flow diagram depicting an exemplary method 900 for limiting a data-flow rate, while FIG. 7 shows a block diagram of a token-bucket queue 710 for metering a data flow based on tokens 732, 733, and 734 deposited therein. Initially, the method 900 involves accepting a first deposit of tokens (e.g., tokens 732, 733, and 734) from the scheduler 401, as indicated at block 910. Typically, the amount of tokens provided within the first deposit is based on, at least one of, the following: the tokens 732, 733, and 734 currently maintained within a token-bucket queue 710, network bandwidth allocated to the subject VM 230 associated with the token-bucket queue 710, or network bandwidth allocated to other VMs located on a common physical node 400 as the subject VM 230.

The method 900 may further involve the steps of consuming the tokens 732, 733, and 734 according to a rate of the data flow from the subject VM 230 (see block 920), and dynamically metering the data-flow rate upon substantially depleting the tokens 732, 733, and 734 (see block 930). In instances of the present invention, the process of metering may include either forwarding to the network interface component 310 packets within the data flow that are received prior to the tokens 732, 733, and 734 being substantially depleted from the token-bucket queue 710, as indicated by block 940. Or, when the tokens 732, 733, and 734 have been substantially depleted from the token-bucket queue 710, the subsequent packets within the data flow may be enqueued, as indicated by block 950.

By way of example, the token-bucket queue 710 may include a token component 730 that periodically receives an amount of tokens (e.g., tokens 732, 733, and 734) deposited thereto by the scheduler 401, and consumes them according to a rate of the data flow from the VM 230. Specifically, as shown in FIG. 7, when the VM 230 transmits the data flow to the token-bucket queue 710, a prescribed amount of tokens 732, 733, and 734 are deducted from the token component 730 prior to releasing the packets 722, 723, and 724 from the queue component 720. The queue component 720 serves to internally enqueue packets that are received by the token-bucket queue 710 when the token component 730 maintains an insufficient quantity of tokens for consumption. As illustrated, packet 721 does not have a corresponding token upon the packets 722, 723, and 724 consuming the currently deposited packages 732, 733, and 734. Accordingly, the queue component 720 will preclude the release of the packet 721 until such time that sufficient tokens are reloaded into the token component 730.

Upon the token component 730 of the token-bucket queue 710 accepting another deposit of tokens from the scheduler 401, the enqueued packet 721 may be released from the queue component 720. In addition, the packet 721 may be forwarded to the network interface component 310, assuming that the tokens within the subsequent deposit are adequate in amount to distribute the packet 721 held by the token-bucket queue 710. This procedure is indicated within the method 900 at block 960.

Turning to FIGS. 4-6, the schemes pertaining to absolute bandwidth and relative weighting will now be discussed. Initially, the “absolute-bandwidth scheme” for distributing network resources involves allocating a particular amount of tokens to a VM regardless of the network resources presently being utilized by other VMs occupying the same physical node. For instance, if the physical node 400 has 10 MBs of available network resources, the VM 270 may be allocated 4 MBs as an absolute value. As discussed above, allocation of 4 MBs may involve depositing an amount of tokens into the token-bucket queue 271 that is equivalent to 4 MBs. If one of the other VMs 230, 235, or 275 assumes an idle state (ceasing to use network resources), the VM 270 will continue to be allocated the consistent 4 MBs of network resources.

In other embodiments, the “relative-weighting scheme” for distributing network resources involves allocating an amount of tokens to a VM that is relative to the network resources presently being utilized by other VMs occupying the same physical node. For instance, referring to FIG. 5, if the physical node 400 has 10 MBs of available network resources, the VM 275 may be allocated 3 MBs as a relative value, where the weighting is  $\frac{1}{3}$  of available bandwidth. As shown, the remaining active VMs 230 and 235 are also provided with a weighting of  $\frac{1}{3}$  (i.e., predefined ratio) of available bandwidth. The allocation of the relative value of 3 MBs involves depositing an amount of tokens into the token-bucket queues 231, 236, and 276 that is equivalent to 3 MBs for each.

As depicted, the VM 270 has assumed an idle state (ceasing to use network resources). Thus, according to the relative-weighting scheme, the network resources initially allocated to the VM 270 are divided amongst the VMs 230, 235, and 237 in the active state. However, in embodiments, the QoS policies may reserve a particular amount of bandwidth for VMs within the idle state. This reserved bandwidth cannot be allocated to VMs in the active state. In the scenario depicted in FIG. 5, 1 MB is reserved for VM 270; thus, the allocated network resources can only combine to be

13

9 MB. It should be noted that, when setting the QoS policies, the total reserve for all VMs on the physical node **400** cannot be greater than the available bandwidth provisioned thereto.

When distributing the tokens, the scheduler **401** may be periodically invoked (e.g., 20 ms), or dynamically triggered upon the occurrence of an event, to reload the token-bucket queues assigned to the VMs that remain in the active condition. The scheduler **401** may detect those VMs that have assumed the idle condition based on the unused tokens residing in their assigned token-bucket queues. As such, the scheduler **401** may abstain from reloading the token-bucket queues of the VMs in the idle state, but may reserve tokens based on their QoS policies, as discussed above.

Further, the scheduler **401** may cap the network resources, or amount of tokens, that may be allocated to a particular VM based on the instructions of the QoS policies. For instance, with reference to FIG. 6, VM **235**, although it is granted a relative value of  $\frac{1}{3}$  of available bandwidth, is capped at 3 MB. Thus, in operation, the cap provides the maximum amount of network resources a data flow can consume, even when a portion of the available bandwidth remains unallocated (e.g., due to idling VMs).

Turning now to FIG. 8 is a flow diagram is illustrated that shows an exemplary method **800** for metering various data flows of packets originating from a plurality of VMs instantiated on the single physical node, in accordance with an embodiment of the present invention. Initially, the method **800** includes the step of inspecting one or more QoS policies written to a contract, as indicated at block **810**. Typically, the contract governs allocation of network resources to a client. Next, tokens may be deposited into a token-bucket queue based, in part, on the QoS policies of the client contract, as indicated at block **820**. As shown in FIGS. 4-6, the token-bucket queue is assigned to the VM instantiated on the physical node. The VM may be identified as having assumed an active state based on a depletion of the deposited tokens from the token-bucket queue, as indicated at block **830**. (When in the active state, a portion of the deposited tokens are withdrawn from the token-bucket queue as a function of a rate of the data flow from the VM.) The method **800** may further include the step of reloading tokens into the token-bucket queue as a function of the QoS policies, a portion of the deposited tokens remaining in the token-bucket queue, or other criteria, as indicated at block **840**.

Embodiments of the present invention have been described in relation to particular embodiments, which are intended in all respects to be illustrative rather than restrictive. Alternative embodiments will become apparent to those of ordinary skill in the art to which embodiments of the present invention pertain without departing from its scope.

From the foregoing, it will be seen that this invention is one well adapted to attain all the ends and objects set forth above, together with other advantages which are obvious and inherent to the system and method. It will be understood that certain features and sub-combinations are of utility and may be employed without reference to other features and sub-combinations. This is contemplated by and is within the scope of the claims.

What is claimed is:

1. One or more computer storage memory having computer-executable instructions embodied thereon that, when executed, perform a method for metering various data flows of packets originating from a plurality of virtual machines (VMs) instantiated on a single physical node, the method comprising:

14

reading one or more first policies contained within a first service agreement, wherein the contract first service agreement governs allocation of network resources to a first client;

allocating a first amount of network bandwidth to a first VM, of the plurality of VMs, associated with the first client based on the one or more first policies of the first service agreement, wherein a size of the first amount of network bandwidth is based, at least in part, upon a first weighting associated with the first VM relative to, at least, a second weighting associated with a second VM of the plurality of VMs, and wherein the first weighting and the second weighting are specified by the first service agreement and a second service agreement, respectively;

detecting that the first VM is in an active state based on a depletion of the first amount of network bandwidth, wherein the first amount of network bandwidth diminishes as a function of a rate of the data flow from the first VM; and

allocating a second amount of network bandwidth to the first VM based, at least in part, on the one or more first policies and a remaining portion of the first amount of network bandwidth.

2. The one or more computer-storage memory of claim 1, the method further comprising:

reading one or more second policies contained within the second service agreement, wherein the second service agreement governs allocation of network resources to a second client;

allocating a third amount of network bandwidth to the second VM, of the plurality of VMs, associated with the second client based on the one or more second policies of the second service agreement;

detecting that the second VM is in an active state based on a depletion of the third amount of network bandwidth, wherein the third amount of network bandwidth diminishes as a function of a rate of the data flow from the second VM; and

allocating a fourth amount of network bandwidth to the second VM based, at least in part, on the one or more second policies and a remaining portion of the third amount of network bandwidth.

3. A computer system capable of metering a rate at which packets from a virtual machine are injected into a network utilizing one or more policies specified by tenants of a data center, the computer system comprising a computer storage medium having a plurality of computer software components embodied thereon, the computer software components comprising:

a first VM instantiated on a physical node, the first VM associated with a first client;

a second VM instantiated on the physical node, the second VM associated with a second client;

a first queue associated with the first VM;

a second queue associated with the second VM; and

a scheduler that reads a first service agreement and a second service agreement associated with the first client and the second client, respectively, that allocates a first amount of bandwidth to the first VM as a function of one or more first policies contained within the first service agreement, and that allocates a second amount of bandwidth to the second VM as a function of one or more second policies within the second service agreement,

15

wherein the first VM transmits a first data flow of packets to the first queue prior to the packets of the first data flow being released from the physical node,

wherein the second VM transmits a second data flow of packets to the second queue prior to the packets of the second data flow being released from the physical node,

wherein, upon receiving the first data flow, the first queue reduces the first amount of bandwidth by a prescribed amount for each packet within the first data flow,

wherein, upon receiving the second data flow, the second queue reduces the second amount of bandwidth by a prescribed amount for each packet within the second data flow.

4. The computer system of claim 3, further comprising: a routing component that accesses a mapping between the first VM and the first queue in response to detecting that the first VM is attempting to transmit the first data flow to an endpoint external to the physical node, and that routes the first data flow to the first queue based on the mapping, wherein the routing component routes the first data flow from first VM to the second VM, bypassing the first queue, in response to detecting that the first VM is attempting to transmit the first data flow directly to the second VM.

5. The computer system of claim 4, the computer software components further comprising a network interface controller (NIC) queue that meters the transmission of the packets within the first data flow from the first queue and, in conjunction, meters the transmission of the packets within the second data flow from the second queue.

6. The computer system of claim 5, the computer software components further comprising a network interface component that receives the metered transmission of the packets within the first data flow and the metered transmission of the packets within the second data flow from the NIC queue and releases from the physical node a controlled volume of traffic that includes the packets within the first data flow and the packets within the second data flow.

7. The computer system of claim 3, wherein the first queue includes a queue component that internally enqueues a subset of packets contained within the first data flow that are received thereby in response to detecting that the first amount of bandwidth is insufficient for consumption.

8. A computerized method for imposing a limit on a rate of data flowing from a subject virtual machine (VM), the method comprising:

accepting a current allocation of bandwidth from a scheduler, wherein an amount of bandwidth provided within the current allocation is based at least in part on a

16

service agreement associated with the subject VM and whether the subject VM is in an active state or an idle state, and

wherein determining whether the subject VM is in an active state is based on a depletion of a previous allocation of bandwidth and determining whether the subject VM is in an idle state is based on a lack of usage of the previous allocation of bandwidth,

dynamically metering a data-flow rate of the subject VM, wherein metering includes:

(a) forwarding, to a network interface component, packets transmitted by the subject VM that are received prior to the current allocation of bandwidth being substantially depleted;

(b) enqueueing packets transmitted by the subject VM that are received subsequent to the current allocation of bandwidth being substantially depleted; and

(c) forwarding the enqueued packets to the network interface component upon accepting another allocation of bandwidth from the scheduler.

9. The computerized method of claim 8, wherein the amount of bandwidth provided within the current allocation is an absolute value specified by the service agreement.

10. The computerized method of claim 8, wherein the amount of bandwidth provided within the current allocation is based further upon a first weighting associated with the subject VM relative to, at least, a second weighting associated with another VM, wherein the another VM is operating on a same physical node as the subject VM.

11. The computerized method of claim 10, wherein the first weighting is specified by the service agreement associated with the subject VM and the second weighting is specified by another service agreement associated with the another VM.

12. The computerized method of claim 8, wherein the amount of bandwidth provided within the current allocation is based further upon an amount of network bandwidth currently available to a physical node on which the subject VM is operating.

13. The computerized method of claim 8, wherein the current allocation is consumed as a function of the data flow rate of the subject VM.

14. The computerized method of claim 8, wherein dynamically metering the data-flow rate of the subject VM further includes:

forwarding, to the network interface component, packets transmitted by the subject VM that are being directly transmitted to another VM operating on a same physical node as the subject VM, regardless of whether the current allocation of bandwidth is substantially depleted.

\* \* \* \* \*